

# Classification using log Gaussian Cox processes

**Alexander G. de. G Matthews**

University of Cambridge  
am554@cam.ac.uk

**Zoubin Ghahramani**

University of Cambridge  
zoubin@eng.cam.ac.uk

## Abstract

McCullagh and Yang (2006) suggest a family of classification algorithms based on Cox processes. We further investigate the log Gaussian variant which has a number of appealing properties. Conditioned on the covariates, the distribution over labels is given by a type of conditional Markov random field. In the supervised case, computation of the predictive probability of a single test point scales linearly with the number of training points and the multiclass generalization is straightforward. We show new links between the supervised method and classical nonparametric methods. We give a detailed analysis of the pairwise graph representable Markov random field, which we use to extend the model to semi-supervised learning problems, and propose an inference method based on graph min-cuts. We give the first experimental analysis on supervised and semi-supervised datasets and show good empirical performance.

## 1 Introduction

Classification is a central problem in statistics and machine learning. Despite the simplicity of the problem specification, and a literature with an age comparable to the fields themselves, it is still an active and diverse research area. This can be explained in part because of the vast variety of applications, which only grows with time.

In this paper we analyze a log Gaussian Cox process (LGCP) model for classification. This model was originally proposed by McCullagh and Yang (2006) as a part of a family of Cox process based models. It has received comparatively little attention in the machine learning and artificial intelligence communities despite being of particular interest when scalable, tractable inference is important.

Perhaps surprisingly, conditioned on the sample covariates, the model reduces to a Potts distribution (Wu, 1982; McCullagh, 2008) for the labels, with the energies simply related to the mean and covariance function of the Gaussian process. There are thus strong connections to familiar models such as the Boltzmann machine (Ackley et al., 1985) and the model can be viewed as a prescribed case of a conditional random field (Lafferty et al., 2001). It is possible to leverage the wealth of literature on inference in such models for both supervised and semi-supervised learning in the LGCP, as we shall show.

The model may be interpreted as a Bayesian nonparametric one (Orbanz and Teh, 2010). Specifying a model in such a way is desirable because it avoids overfitting (i.e Bayesian) and has large support (i.e nonparametric). The particular model in this case is a marked Cox process (Cox, 1955), which is a particular type of spatial point process (Daley and Vere-Jones, 2003). The latent intensity for each class is modelled using an exponentiated Gaussian process (Møller et al., 1998).

In the supervised case there are intriguing relations to classical nonparametric methods such as kernel density estimation (Rosenblatt, 1956; Parzen, 1962) and nearest neighbour classifiers, raising interesting questions about the interplay between classical nonparametric statistics and Bayesian nonparametric methods.

It is desirable that inference in a model be simple and robust. Scaling to large datasets can often be challenging with Bayesian nonparametric methods. For example, the intrinsic computational complexity of Gaussian process regression is  $O(N^3)$  (Rasmussen and Williams, 2006) where  $N$  is the number of training points. Further, Gaussian process classification requires approximate inference. By contrast, the supervised LGCP model has  $O(N)$  closed form computation for prediction on a test point which means that we have been able to investigate our model on some large datasets. So far fairly simple parameter training using cross validation has been sufficient, though it seems likely this could be developed using the literature on Markov random fields (Koller and Friedman, 2009).

## 1.1 Relation to existing work and contributions of this paper

As we have already described, this work builds on the excellent contribution of McCullagh and Yang (2006), which has been discussed in subsequent literature.

In the first paper of McCullagh and Yang (2006), the LGCP classifier and the related permanental Cox process model are first considered. The supervised model equations, including the predictive equations (17) are motivated for the first time but the LGCP model is not implemented or tested empirically. A connection to Gibbs partitions (Pitman, 2006) is mentioned.

In a paper on bias in logistic models (McCullagh, 2008), the two models are used as an illustrative example of a non-standard supervised classification model. The relation between the Janossy density and product density is discussed in the context of different sampling schemes and a new derivation (which we have followed) is given. The connections to Gaussian process classification, here seen as a special case of logistic regression with random effects, are discussed. There is a single sentence describing the connection to Markov random fields (Besag, 1974).

In a paper on the permanental Cox process classifier (Yang et al., 2012), which could be viewed as an analogue of this paper, the permanental variant is implemented using an approximation scheme based on a cyclic approximation, and then tested on real data. For further discussion of the relation between the permanental model and the log Gaussian model see (McCullagh and Yang, 2006; McCullagh, 2008) and for some summary comments given for the perspective of machine learning see the appendix.

In the field of semi-supervised learning, Markov random fields and the use of graph min-cuts for MAP inference have been considered before in the literature (Zhang et al., 2001; Blum and Chawla, 2001) but not with a Bayesian nonparametric motivation. As shall be seen in section 3.4 the Bayesian nonparametric model places specific requirements on the edge weights of the Markov random field that are prescribed by the prior.

To summarize, the contributions of our work are, in broad terms:

- The recognition of the LGCP classification model from McCullagh and Yang (2006), which despite a variety of appealing characteristics has not, to our knowledge, been previously used in practice.
- The discovery of new connections between the supervised LGCP method and classical nonparametric methods as discussed in section 3.2.1. These additional connections prove important in explaining the experimental results.
- The first empirical validation of the supervised method against some commonly used datasets in section 4.2, including some large, high dimensional datasets. Large dataset results are of particular interest in Bayesian nonparametrics, where they can be harder to obtain.
- To elaborate, in section 3.3, the general form of the crucial link to Markov random fields, after the initial sentence from McCullagh (2008).
- To show, in section 3.5, that the Markov random field is pairwise graph-representable in the sense of Kolmogorov and Zabih (2004) and hence pairwise submodular.
- To extend the model to semi-supervised learning in section 3.4.
- To propose robust inference for the semi-supervised learning model using graph min-cut algorithms, that exploit the pairwise graph-representable property, as described in section 3.4.
- The experimental validation of the proposed semi-supervised learning algorithm on real data as described in section 4.3.

In this paper we now proceed by providing a summary of the necessary background theory, which we have strived to make relatively self-contained.

## 2 Background

We will be interested in this paper in spatial point processes (Daley and Vere-Jones, 2003; Møller and Waagepetersen, 2004). The realisations of spatial point processes are sets  $X$  of points  $x$  in some space. In this paper all the points will

lie in  $\mathbb{R}^D$ . Each point will be supplemented with a mark, which in this case will be in some finite label set  $\mathcal{L}$ . The pairs of points and labels constitute what is called a marked point process.

The well known Poisson process (Kingman, 1993) is a simple spatial point process that can be used to build more complex models. It is parameterized by a non-random intensity function  $\rho : \mathbb{R}^D \mapsto [0, \infty)$  which has the property of being locally integrable, so that  $\int_A \rho(x)dx < \infty$  for all bounded Borel sets  $A \subset \mathbb{R}^D$ . This can be used to define an intensity measure so that

$$\eta(A) = \int_A \rho(x)dx \quad (1)$$

This measure is assumed to be diffuse, which means that  $\eta(\{x\}) = 0 \ \forall x \in \mathbb{R}^D$ . The Poisson process can then be defined in terms of the distribution over the number of points in any given Borel set  $A$ , which we denote as  $n(A)$ .

1. For Borel set  $A \subset \mathbb{R}^D : \eta(A) < \infty$ , the distribution of  $n(A)$  is Poisson, with rate parameter  $\eta(A)$  if  $\eta(A) \neq 0$  and  $n(A) = 0$  otherwise.
2. For Borel sets  $A, B$  such that  $A \cap B = \emptyset$ ,  $n(A)$  is independent of  $n(B)$ .

The Cox process (Cox, 1955) can be defined in terms of the Poisson process. Instead of a non-random intensity  $\rho$ , we now make it a well-behaved non-negative random function. Then conditioned on this random function, the points from a Cox process are distributed as a Poisson process with that intensity.

The case where the random function is given by the exponential of a Gaussian process is called a log Gaussian Cox process (Møller et al., 1998).

$$f \sim \mathcal{GP}(\mu, C) \quad (2)$$

$$\rho = \exp(f) \quad (3)$$

Here  $\mathcal{GP}$  denotes a Gaussian process (Rasmussen and Williams, 2006). Following the exposition of Møller et al (1998), some restrictions on the form of this Gaussian process are assumed in this paper. First it is assumed that the process is stationary, which will mean that the mean function  $\mu$  is constant and that the covariance function  $C(x_j, x_k)$  has the form of a function  $C(x_j - x_k)$ . Further it is assumed that there exist  $\alpha, \beta > 0$  such that:

$$1 - \frac{C(s)}{C(0)} < \alpha \|s\|^\beta \ \forall s : \|s\| < 1 \quad (4)$$

This last condition ensures the intensity measure  $\eta$  is almost surely a continuous modification of  $f$ . These are sufficient, but not necessary conditions for a well defined LGCP. For further detail on such matters see the paper of Møller et al (1998). Examples of valid kernels include the commonly used squared exponential, and exponential kernels.

Figure (1) demonstrates the LGCP in two dimensions. The Gaussian process had mean zero and a squared exponential kernel, with unit length scale and signal variance. A realisation of the random intensity of the Cox process and corresponding sample points are shown.

An important set of quantities for describing a Cox process, if they exist is the product densities (Daley and Vere-Jones, 2003). The  $K$ -th product density, where  $K$  is a positive integer, is a function  $m_{[K]}(x_1, x_2, \dots, x_K)$  of points  $x_j \in \mathbb{R}^D$ , for each  $j = 1, \dots, K$ . An intuitive interpretation of the product density can be given via the equation:

$$\begin{aligned} m_{[K]}(x_1, x_2, \dots, x_K) dx_1 dx_2 \dots dx_K \\ = Pr(\text{One point in each small volume } (x_j, x_j + dx_j)) \end{aligned} \quad (5)$$

This can be contrasted to the  $K$ -th local Janossy density  $\theta_K(x_1, x_2, \dots, x_K | A)$ , defined relative to some bounded Borel set  $A \subset \mathbb{R}^D$  which, when it exists, has the property:

$$\begin{aligned} \theta_K(x_1, x_2, \dots, x_K | A) dx_1 dx_2 \dots dx_K \\ = Pr(\text{Exactly } K \text{ points in } A, \text{ one in each small volume } (x_j, x_j + dx_j) \text{ and none elsewhere in } A) \end{aligned} \quad (6)$$

Clearly, the two differ in the sense that the local Janossy density quantifies occurrence of points and non-occurrence of points elsewhere, whereas the product density only quantifies the occurrence of points in a certain region, and does




Figure 1: Left: sample of random intensity function for a log Gaussian Cox process with isotropic squared exponential kernel with unit variance and length scale. Right: sample points drawn with a Poisson rate from the same intensity function (see text)

not rule out the possibility of points elsewhere. The local Janossy density is the density of the local Janossy measure and the product density is the density of the factorial moment measure, both with respect to Lebesgue measure on  $(\mathbb{R}^D)^K$ . For more details see Daley and Vere-Jones (2003).

In the case of the LGCP, subject to conditions that are always met in this paper, the product densities exist (Møller et al., 1998) and are given by:

$$m_{[K]}(x_1, x_2, \dots, x_K) = \exp \left\{ K\mu + \frac{1}{2} \sum_{j,k=1}^K C(x_j - x_k) \right\} \quad (7)$$

For what follows we shall, for convenience, refer to the product densities and Janossy densities in the singular since the dependence on  $K$  is obvious from the context.

## 3 Model

### 3.1 Model definition

In this section we review the model for LGCP classification given in the original paper by McCullagh and Yang (2006), although the treatment of the observation model follows the later work of McCullagh (2008).

A fixed, known, number of classes  $Q$  is assumed. Each class population  $i = 1, \dots, Q$  is modelled using a separate independent log Gaussian Cox process  $X_i$  with mean function  $\mu^{(i)}$  and covariance function  $C^{(i)}$ .

$$f_i \sim \mathcal{GP}(\mu^{(i)}, C^{(i)}) \quad i = 1, \dots, Q \quad (8)$$

$$\rho_i = \exp(f_i) \quad (9)$$

$$X_i \mid \rho_i \sim \mathcal{PP}(\rho_i) \quad (10)$$

Here  $\mathcal{PP}(\cdot)$  denotes a Poisson process with a given intensity. If one considers the superposition of the  $Q$  different Cox processes and ignores the labels, another Cox process  $X$  is obtained, which is termed the superposition process.

$$X = \cup_{i=1}^Q X_i \quad (11)$$

This Cox process may be associated with the total intensity  $\rho_T = \sum_{i=1}^Q \rho_i$ . The superposition Cox process is not a *log Gaussian* Cox process since the distribution for  $\rho_T$  cannot be found by exponentiating any Gaussian process.

Next a method is defined by which samples from the population come to be observed. Let  $\tau(S, \gamma)$  be the spatial point process whose points lie in  $\mathbb{R}^D$  resulting from an independent random thinning of the points in  $S$ . Loosely speaking this means that with probability  $\gamma$  a point in  $S$  is present in the new set of points. It is assumed that this thinning procedure is applied to each of the class populations to obtain new thinned category samples  $T_i$ :

$$T_i | X_i \sim \tau(X_i, \gamma) \quad (12)$$

One may also define the thinned superposition  $T = \cup_{i=1}^Q T_i$ . Consider fixing  $A$  to be the bounded Borel set in which observations are performed and let  $\mathbf{x} = A \cap T$  be the observed covariates. Now define  $\mathbf{x}^{(i)} = \{x_j \in \mathbf{x} : y_j = i\}$ , which denotes the set of points in the sample that have label  $i$ . The conditional probability, for a given value of  $\gamma$  is equal to:

$$Pr_\gamma(\mathbf{y}|\mathbf{x}) = \frac{\mathbb{E} \left[ \exp\{-\gamma \int_A \rho_T(x') dx'\} \prod_{j=1}^N \rho_{y_j}(x_j) \right]_{\{\rho_i\}}}{\mathbb{E} \left[ \exp\{-\gamma \int_A \rho_T(x') dx'\} \prod_{j=1}^N \rho_T(x_j) \right]_{\{\rho_i\}}} \quad (13)$$

where the expectation is taken over the set of intensity functions  $\{\rho_i\}_{i=1}^Q$ . Consider the limit  $\gamma \rightarrow 0$  which represents an idealized observation process. In the limit the distribution of the labels is given by:

$$Pr(\mathbf{y})[\mathbf{x}] = \lim_{\gamma \rightarrow 0} [Pr_\gamma(\mathbf{y}|\mathbf{x})] = \frac{\prod_{i=1}^Q m^{(i)}(\mathbf{x}^{(i)})}{M(\mathbf{x})} \quad (14)$$

$m^{(i)}(\cdot)$  is the product density of Cox process  $i$  and is given by equation (7) evaluated with the corresponding mean function  $\mu^{(i)}$  and covariance function  $C^{(i)}$ .  $M(\cdot)$  is the product density of the superposition process, which, as we have discussed, does not have the form of the product density for an LGCP. We will investigate its form in the sections that follow. It would be tempting to denote this distribution using the conditional notation ' $Pr(\mathbf{y}|\mathbf{x})$ ' but we use the non-standard notation  $Pr(\mathbf{y})[\mathbf{x}]$  to highlight the limiting construction being used.

Post-hoc it can be seen that the label distribution in equation (14) has an intuitive interpretation in the sense that it is proportional to a product of product densities, one per class divided by the superposition product density. Using the interpretation given in equation (5) one sees the relationship to the probability that each independent process contains a certain set of points in a small interval around  $\mathbf{x}^{(i)}$  in precisely the sense of a product density as opposed to a Janossy density. Note also that the dependence on the observation Borel set  $A$  has vanished in the limit. An advantage of this observation model, which at least partly motivates the approach (McCullagh, 2008) is the relative tractability of the result.

Substituting the product density from equation (7) into equation (14) one obtains:

$$Pr(\mathbf{y})[\mathbf{x}] = \frac{1}{M(\mathbf{x})} \exp \left\{ \sum_{i=1}^Q |\mathbf{x}^{(i)}| \mu^{(i)} + \frac{1}{2} \sum_{i=1}^Q \sum_{(x_j, x_k) \in \mathbf{x}^{(i)} \times \mathbf{x}^{(i)}} C^{(i)}(x_j - x_k) \right\} \quad (15)$$

where  $|\mathbf{x}^{(i)}|$  denotes the cardinality of the set  $\mathbf{x}^{(i)}$ .

Cox (1958) and McCullagh (2008) consider the *no interference condition*. In general a discriminative model fulfils the no-interference condition if it has the following property:

$$\sum_{\mathbf{y}^*} Pr(\mathbf{y}^*, \mathbf{y}|\mathbf{x}^*, \mathbf{x}) = Pr(\mathbf{y}|\mathbf{x}) \quad (16)$$

For example, the multivariate Gaussian marginalization property guarantees that this criterion is fulfilled in the case of Gaussian process classification (Rasmussen and Williams, 2006). This means that knowledge of the existence of other unlabelled covariates does not affect our predictions about data. McCullagh (2008) discusses the extent to which this is a desirable characteristic of a model. Here we point out that for the concept of semi-supervised learning to make sense in a model the no interference condition must not be fulfilled. Treating the limiting process as standard conditioning the LGCP model does not fulfil the no-interference condition and thus has a meaningful extension to semi-supervised learning as we discuss in section 3.4.

## 3.2 Supervised learning

The supervised learning case was originally discussed by McCullagh and Yang (2006). Here we review the existing material and discuss in section 3.2.1 some novel connections to classical nonparametric methods. These new connections prove important for understanding the behaviour of the supervised method which will be evaluated empirically for the first time in the experiments section.

Consider the case where a set of covariates  $\mathbf{x}$  and their labels  $\mathbf{y}$  have been observed and one wishes to predict the label  $y^*$  of a test point with measured covariates  $x^*$ . The requisite distribution is:

$$Pr(y^* = i | \mathbf{y})[\mathbf{x} \cup \{x^*\}] = \mathcal{S}(F)_i \quad (17)$$

We have defined the softmax function  $\mathcal{S} \in \mathcal{R}^Q$  which has components:

$$\mathcal{S}(x)_i = \exp(x_i) / \left( \sum_{\beta=1}^Q \exp(x_\beta) \right) \quad (18)$$

and  $F \in \mathbb{R}^Q$  whose components are given by:

$$F_i = \mu^{(i)} + \frac{1}{2} C^{(i)}(x^*, x^*) + \sum_{x \in \mathbf{x}^{(i)}} C^{(i)}(x^*, x) \quad (19)$$

In the case where the mean and covariance are the same for each category the first two terms in  $F_i$  will cancel when one takes the softmax. Clearly computation of the predictive density is linear in the number of training data points.

### 3.2.1 Connections to classical nonparametric methods.

We now show a novel relation of this model to *kernel density* or *Parzen-Rosenblatt window* (Rosenblatt, 1956; Parzen, 1962) estimates of the individual class densities. These new connections prove important for understanding the behaviour of the supervised method which will be evaluated empirically for the first time in the experiments section.

Given a symmetric non-negative normalized function of two points  $G(x, x')$  the simplest kernel density estimate  $\beta(x^*)$  at a test point  $x^*$  given a set of  $N$  points  $x_j$  is given by:

$$\beta(x^*) = \frac{1}{N} \sum_{j=1}^N G(x^*, x_j) \quad (20)$$

Now let us assume that we take the kernel density estimate of each class  $i$  each containing  $N_i$  points and giving  $N_T$  points in total. We wish to produce a probabilistic estimate of the class of  $y^*$ . Allowing for the empirical counts in each class with a factor  $\frac{N_i}{N_T}$  we obtain:

$$Pr(y^* = i) = \frac{\beta(x^*)_i N_i}{\sum_{j=1}^Q \beta(x^*)_j N_j} \quad (21)$$

If we stipulate that:

$$G(x, x') = C^{(i)}(x, x') = C(x, x') \quad \forall i \quad (22)$$

and that  $\mu^{(i)} = 0 \quad \forall i$  then the predictive equations (17) and (21) imply some functional similarities and differences between the two models. Clearly the predictive probability distributions are not the same since the LGCP has a softmax function. However under a 0-1 loss the optimal decision rule for classification is to choose the maximum assignment. Both predictive distributions in fact have the same maximum leading to the same decisions. This equivalence extends to the parameter estimation if we use cross-validation on the 0-1 loss but not if we for instance use the negative log of the predictive probability. In the case where the covariance functions are different it would seem hard to ‘derive’ the term  $\frac{1}{2} C^{(i)}(x^*, x^*)$  in equation (19) in terms of the kernel density estimator.

The kernel density estimator is well studied in nonparametric statistics. In the regimes where they are equivalent, frequentist guarantees derived for the kernel density style algorithms apply equally to LGCP supervised learning. In the regimes where they are not equivalent this would seem like a good starting point to derive such guarantees.

### 3.3 Relation to conditional Markov random fields

Equation (14) is related to the many types of Markov random field (Besag, 1974), as is briefly mentioned in a sentence in a paper of McCullagh (2008). In this section then sections 3.4 and 3.5 we further analyze this link, giving the first detailed exposition. Our augmented characterization will prove essential for the generalization to semi-supervised learning given in section 3.4.

Equation (15) may be rearranged to give:

$$Pr(\mathbf{y})[\mathbf{x}] = \frac{1}{M(\mathbf{x})} \exp \left\{ \sum_{j=1}^N \mu^{(y_j)} + \frac{1}{2} \sum_{j,k=1}^N C^{(y_j)}(x_j - x_k) \delta(y_j, y_k) \right\} \quad (23)$$

where  $\delta$  denotes the Kronecker delta function. The connection to Markov random fields has thus been made explicit. This probability distribution is a type of fully connected Potts model (Wu, 1982). In general the pairwise energies depend not only on the covariates  $\{x_j\}$  but also on the classes  $\{y_j\}$ . The product density of the superposition process  $M(\mathbf{x})$ , can be seen to correspond to the partition function or normalizing constant of the Potts model.

The link to the Potts model here differs from their use in some other pattern recognition problems, in the sense that the weights are stipulated by the kernel function of the Gaussian process, which are in turn derived from the observed covariates. In this sense, conditioned on  $\mathbf{x}$ , one can also view the resulting Potts model as a kernelized conditional random field (Lafferty et al., 2001).

If we choose to make the simplifying assumption that the means  $\mu^{(i)}$  are zero and the kernel functions are all the same  $C^{(i)} = C$ . The conditional model can be written as:

$$Pr(\mathbf{y})[\mathbf{x}] = \frac{1}{M(\mathbf{x})} \exp \left\{ \frac{1}{2} \sum_{j,k=1}^N C(x_j - x_k) \delta(y_j, y_k) \right\} \quad (24)$$

If we take another special case, with two categories ( $Q = 2$ ), then we obtain a fully connected Ising model or Boltzmann machine (Ackley et al., 1985).

In general, both the Ising model and the Potts model partition functions are  $\#P$ -Hard to compute, even with the constraint of positive weights. Thus, so is the product density of the superposition process. However, in the positive weight, binary case, it is possible to approximate these partition functions with high probability, to within a given percentage, in polynomial time using Markov chain Monte Carlo methods (Jerrum and Sinclair, 1993).

We delay our discussion of the pairwise graph-representability and submodularity of the Markov random field until section 3.5.

### 3.4 Semi-supervised learning

In this section, using the connection to the Potts model, we show how to extend the model to semi-supervised learning for the first time, then discuss the relationship to ideas in the existing semi-supervised learning literature.

Consider a partitioning of the data into two sets; the ‘starred’ set  $\mathbf{x}^*, \mathbf{y}^*$  which will correspond to the unlabelled test data and the ‘unstarred set’  $\mathbf{x}, \mathbf{y}$  which will correspond to the labelled data. Conditioned on the training covariates  $\mathbf{x}$  and a set of test covariates  $\mathbf{x}^*$ , the semi-supervised learning problem takes the form of inference for a Boltzmann machine with hidden units  $\mathbf{y}^*$  corresponding to the set of unknown labels. The quantity of interest is:

$$Pr(\mathbf{y}^*|\mathbf{y})[\mathbf{x}^* \cup \mathbf{x}] \quad (25)$$

This obeys the relation:

$$Pr(\mathbf{y}^*|\mathbf{y})[\mathbf{x} \cup \mathbf{x}^*] = \frac{Pr(\mathbf{y}^*, \mathbf{y})[\mathbf{x} \cup \mathbf{x}^*]}{Pr(\mathbf{y})[\mathbf{x} \cup \mathbf{x}^*]} \propto Pr(\mathbf{y}^*, \mathbf{y})[\mathbf{x} \cup \mathbf{x}^*] \quad (26)$$

Here proportionality  $\propto$  means the conditional distribution for  $Pr(\mathbf{y}^*|\mathbf{y})[\mathbf{x} \cup \mathbf{x}^*]$  will have the same functional form for  $\mathbf{y}^*$  as  $Pr(\mathbf{y}^*, \mathbf{y})[\mathbf{x} \cup \mathbf{x}^*]$  up to a multiplicative constant. This constant of proportionality will not affect the MAP assignment for  $\mathbf{y}^*$  which we will target for semi-supervised learning.

We define an ‘energy’  $E(\mathbf{y}^*)$  that obeys:

$$Pr(\mathbf{y}^*, \mathbf{y})[\mathbf{x} \cup \mathbf{x}^*] = \exp\{-E(\mathbf{y}^*)\} \quad (27)$$

The MAP estimate for  $\mathbf{y}^*$  corresponds to the minimum energy configuration. Following the convention of Boykov et al (2001), we consider the energy to be a sum of three terms- a ‘data’ term, a ‘smoothing’ term, and a constant term  $c$  that does not depend on  $\mathbf{y}^*$ .

$$E(\mathbf{y}^*) = E_{\text{data}}(\mathbf{y}^*) + E_{\text{smooth}}(\mathbf{y}^*) + c \quad (28)$$

To ascertain the first two terms we compare to equation (23). The terms are given by:

$$-E_{\text{data}}(\mathbf{y}^*) = \sum_{y_i^*}^{|x^*|} \mu^{(y_i^*)} + \frac{1}{2} \sum_{j=1}^{|x^*|} C^{(y_j^*)}(\mathbf{0}) + \sum_{j=1}^{|x^*|} \sum_{k=1}^{|x^*|} \delta(y_j, y_k^*) C^{(y_j)}(x_j, x_k^*) \quad (29)$$

and:

$$-E_{\text{smooth}}(\mathbf{y}^*) = \sum_{j < k=1}^{|x^*|} \delta(y_j^*, y_k^*) C^{(y_j^*)}(x_j^* - x_k^*) \quad (30)$$

Note that we no longer have a factor of 1/2 in equation (30) because we now sum only over  $j < k$ . We write it this way for consistency with Boykov et al (2001).

Recall that all kernel functions we consider are non-negative. Under these conditions, as we will show in the next section 3.5, it is possible to exploit graph min-cut methods. In the case of binary classification this will give the exact MAP state in low order polynomial time (Kolmogorov and Zabih, 2004). For more than two states we can no longer guarantee to find the MAP estimate in polynomial time. However we can use the expansion algorithm (Boykov et al., 2001) to find a strong local optimum.

As has already been mentioned, semi-supervised learning using general Markov random fields has been considered before in the machine learning community (Zhang et al., 2001; Blum and Chawla, 2001). The latter of these is the closest to the current work since it uses graph min-cut algorithms. The comparison here is informative. Our work derives the model from a log Gaussian Cox process which is not true of any algorithm in the existing literature. The derivation from a log Gaussian Cox process places additional requirements on the energy functions whilst (as we shall prove) still maintaining the submodularity necessary for efficient inference. As well as being interesting from a Bayesian nonparametric perspective, these extra requirements seem to answer the difficulty that Blum and Chawla describe in motivating and choosing the form of energy function for their algorithm, when viewed as an MRF. One potential issue with using the global MAP assignment is that it is the optimal decision under a global 0-1 loss rather than the Hamming loss. Blum et al discuss these issues (2004) and advocate an ensemble based on randomizing the potential functions. The empirical improvements they observe can perhaps in hindsight be related to recent work on approximately sampling MRFs in the ‘strong data/strong coupling’ regime using randomized MAP methods (Hazan et al., 2013). We expect such considerations are relevant here but they are beyond the scope of the current paper.

### 3.5 Pairwise graph representability and submodularity of Markov random field

In this section we describe new results on the pairwise graph representability and submodularity of the model. This is useful because it means it is possible to use low order polynomial time graph min-cut inference.

Kolmogorov and Zabih (2004) define the term *graph-representable* for a binary energy minimization problem that can be represented in such a way that it is possible to find its minimum using graph min-cut algorithms. Min-cut algorithms find the exact minimum energy configuration of the binary problem. The same work contains the following result:

**Theorem 1** (Kolmogorov and Zabih (2004)). *A real valued energy function  $E(\mathbf{y}^*)$  with binary inputs  $\mathbf{y}^* \in \{0, 1\}^N$  and:*

$$E(\mathbf{y}^*) = \sum_j E_j(y_j^*) + \sum_{j < k=1}^N E_{(j,k)}(y_j^*, y_k^*) \quad (31)$$

*is graph-representable if and only if:*



$$E_{(j,k)}(0,0) + E_{(j,k)}(1,1) \leq E_{(j,k)}(0,1) + E_{(j,k)}(1,0) \quad (32)$$

for all pairs  $(j, k)$ .

For a proof of this result and further discussion see the original work. For more than two classes the expansion algorithm breaks the problem down repeatedly into a binary problem and then applies a graph min-cut algorithm. Although the expansion algorithm is not guaranteed to find the exact minimum energy state it finds a strong local optimum. For more discussion of the quality of this approximation see the thesis of Veksler (1999). The condition for the expansion algorithm to be valid is that the energies satisfy (Kolmogorov and Zabih, 2004):

$$E_{j,k}(a, a) + E_{j,k}(b, c) \leq E_{j,k}(a, c) + E_{j,k}(b, a) \quad (33)$$

for any three classes  $a, b, c$  and sites  $j, k$ . We refer to this condition as *pairwise graph-representability*. It can be seen that pairwise graph-representability corresponds to graph-representability when we are considering the binary case only. We now give a new result.

**Theorem 2.** *The energy minimization problem corresponding to the maximum a posteriori state of the log Gaussian Cox process semi-supervised learning problem is pairwise graph representable.*

To prove this we need to show that condition (33) holds for the energy functions in our problem. It suffices only to consider the pairwise terms given in equation (30). That is we take:

$$E_{j,k}(y_j^*, y_k^*) = E_{j,k}(y_j^*, y_k^*)_{\text{smooth}} = -\delta(y_j^*, y_k^*) C^{(y_j^*)}(x_j^* - x_k^*) \quad (34)$$

We simply need to exhaust each relevant possibility for the relationship between the labels. That is to say we know that:

$$(a = b = c) \vee ((a = b) \wedge (b \neq c)) \vee ((a \neq b) \wedge (b = c)) \vee ((a \neq b) \wedge (b \neq c)) \quad (35)$$

where  $\vee$  denotes exclusive logical disjunction. We then substitute each case into the inequality to verify it is correct. We will need to use the condition that all covariance functions are assumed to be non-negative.

Pairwise graph representable problems are a subset of pairwise submodular problems (Kolmogorov and Zabih, 2004). Hence a corollary of our theorem is that the problem is pairwise submodular. More general submodular solvers are not currently generally as fast as graph min-cut solvers where the latter are applicable since the former cannot exploit the extra specific structure present in the problem.

### 3.6 Parameters

The effect of multiplying mean and covariance functions by some common value  $\gamma > 1$  is analogous to changing the temperature in a Markov random field. It will increase the confidence of predictions in equation (17) but it will not change the order of the probabilities and hence not change the most likely category.

Under the stationarity assumption, we have a constant mean function for each category. If we consider the effect on equation (17) of changing the means, we can see that it allows us to model a prior bias between categories, independent of the training data.

## 4 Experiments

To perform experiments we created software in MATLAB, building on the GPML toolbox (Rasmussen and Nickisch, 2010). For the graph min-cut algorithms we used the code that accompanies the original papers (Boykov et al., 2001; Boykov and Kolmogorov, 2004; Kolmogorov and Zabih, 2004). Comparisons to the support vector machine (SVM) were performed using SVM-light (Joachims, 1999). The harmonic functions algorithm we used for comparison in the semi-supervised experiments was a slight modification of the code accompanying the original paper (Zhu et al., 2003).

## 4.1 Experiments on demonstration data

In order to demonstrate the LGCP classifier we first show performance on an illustrative synthetic dataset. Figure (2) (left) shows a generated dataset of two concentric circles. The outer training data is clearly not in a convex set, but this is not an issue for the method. Figure (2) (centre) shows the result of performing prediction on a notional test point using equation (17). The mean functions were taken to be zero and an isotropic squared exponential kernel with signal standard deviation 0.5 and length scale 1 was used. It can clearly be seen that the model gives sensible predictions on this test data. Figure (2) (right) shows a semi-supervised version of the same experiment with the same fixed parameters for the kernel. Using the graph min-cut methods described in section 3.4 the labels are recovered with only one error, made on a point that is close to the margin between classes. We performed the analogous experiment for three concentric circles, again using the graph min-cut algorithm for the semi-supervised learning, and obtained similar good results.




Figure 2: Left: Generated data with two concentric circles used for demonstrating LGCP. Centre: Predictive probability for a notional test point in the 2D plane for the same data. Right: A semi-supervised version of the same experiment. The labelled data are shown in red and blue and the unlabelled data are shown in green. The labels are recovered with only one error circled in green in the left panel.

## 4.2 Experiments on supervised classification datasets

Experiments were performed on six commonly used classification datasets. The Synth dataset, and the Diabetes dataset come from Ripley (Ripley, 1996)<sup>1</sup>. The Banana dataset originally comes from Rätsch<sup>2</sup>. They are commonly used for comparison of classification algorithms, see for instance (Naish-Guzman and Holden, 2008). The Brain Computer Interface (BCI) dataset is originally used in the work of Lal et al (2004). The covariates correspond to the parameters of a time series model fitted to electroencephalography data taken from a single subject. The categories correspond to the subject imagining a movement either with the left or right hand. The restricted STL (rSTL) dataset is a subset of the full STL image dataset (Coates et al., 2011). The dataset consists of  $96 \times 96$  RGB images. We randomly selected 500 images each of monkeys and trucks. Both the BCI dataset and rSTL dataset are relatively challenging classification tasks where the number of data points is in some sense small relative to the dimensionality of the data. The MNIST handwritten digit dataset (LeCun et al., 1998, 2014) is a very highly investigated classification dataset. The data has a relatively large number of data points, is multiclass, and has relatively high dimensional covariates.

For the Synth, Diabetes, and MNIST datasets we used the standard partition to aid comparison to other papers in the literature. For the Banana dataset, BCI and rSTL we randomly partitioned between train and test.

The experiments were performed using the LGCP classifier with zero mean function. All experiments used the isotropic squared exponential kernel, which was shared between classes. In all cases the single shared length scale

<sup>1</sup>Data available from <http://www.stats.ox.ac.uk/pub/PRNN/>

<sup>2</sup>Data available from <http://mldata.org/>

was the only free parameter. It was fitted using leave-out-one cross validation with a 0-1 loss. The parameter training was done using all the training points, with the first five smaller datasets and 1000 randomly chosen points from the MNIST training set. In the latter case this was sufficient for our needs, given that we were only training one parameter. One iteration of cross validation over all 1000 points, for a given length scale, on this MNIST subset took 17 seconds on a single Intel I3 2.53 GHz core with 4GB of RAM. The prediction on the validation data for MNIST using all 60000 training points and 10000 validation points took 130 seconds on the same computer.

For the datasets other than MNIST we compared to the support vector machine (SVM) and Gaussian process classification (GPC) both using the same kernel. For GPC we used variational approximate inference, the logistic link function, and zero mean function. To mitigate any issues with local optima we repeated the GPC gradient descent with 10 different random initializations and took the run with the best marginal likelihood. The kernel length scale of the SVM was trained using leave-out-one cross validation whilst the regularization parameter was set using the default heuristic of the software. For the MNIST dataset a wide variety of algorithms have been tested under controlled conditions and the results published. Insofar as we are aware there is no published result for GPC on the MNIST dataset which is surprising because of the significance of such a result. We speculate that this is due to the technical challenges of providing multiclass, sparse, approximate inference on a large high dimensional dataset.

Table 1: Error \% for supervised binary datasets

Dataset	Train /Test	Variables /Classes	LGCP	GPC	SVM
SYNTH	250 /1000	2 /2	9.1	10.1	9.4
DIABETES	200 /332	7 /2	21.7	20.8	22.3
BANANA	400 /4900	2 /2	11.4	10.6	10.7
RSTL	500 /500	27648 /2	20.0	18.6	12.4
BCI	200 /200	117 /2	39.5	24.5	31.5

Table 2: Error \% for Supervised MNIST Dataset

Train /Test	Var /Classes	LGCP	1-NN	SVM	FF. net	Log. reg.	RBF	Conv. net
60K/10K	784 /10	3.11	3.09	1.4	4.5	12.0	3.6	0.23

The results obtained for the datasets excluding MNIST are shown in table 1. The results for LGCP on the Synth, Diabetes and Banana datasets are comparable with the SVM and GPC. This is the case despite both of these two alternative algorithms having more free parameters than the single parameter LGCP. The LGCP classifier lags the SVM and GPC on the two higher dimensional datasets BCI and rSTL although neither of these alternatives is better outright.

The results for LGCP on the MNIST dataset, shown in table 2, are similar to those published for  $K$ -nearest neighbours (KNN) with a Euclidean  $L_2$ -norm. Given the links to classical nonparametric models described in section 3.2.1 this is perhaps not surprising. For sufficiently short length scales the sum over training points in equation (19) is dominated by the term with the smallest Euclidean distance, i.e the nearest neighbour. Conditions under which kernel sum based algorithms will be similar to KNN are discussed in the literature (Bengio et al., 2005). If the data lie near a low dimensional manifold the squared exponential kernel will only depend on the dimensionality of this manifold (Tenenbaum et al., 2000) which is also of relevance to semi-supervised learning.

The method is invariant to joint permutation of the covariates. Despite its simplicity it outperforms some other more complex methods (LeCun et al., 1998, 2014) such as early experiments with the feed forward neural network (4.5%). The SVM with the same kernel performs well on this dataset with an error of 1.4%. The best current performance is achieved by deep convolutional neural networks, which can achieve very low error rates around 0.23% (Ciresan et al., 2012). These methods are not permutation invariant, incorporate strong prior knowledge, and require a good deal of training.

The LGCP algorithm has scope for generalization since there is a wide variety of mean functions and covariance functions that would encode more complex inductive biases, particularly if as we suspect, the assumption of stationarity may be relaxed. Since this would lead to a larger number of parameters it would be necessary to leverage the existing literature on approximate learning in Markov random fields (Koller and Friedman, 2009).

### 4.3 Semi-supervised experiments

We also performed experiments to compare SVM, GPC, LGCP, semi-supervised LGCP (SLGCP) and the semi-supervised Harmonic functions algorithm of Zhu et al (2003).

To compare the algorithms we used three different datasets. The double helix dataset is a synthetic dataset inspired by the paper of Zhu et al (2003). An advantage of this dataset is that it is a non-trivial classification task where it is still possible to visualize the geometric advantage that semi-supervised methods have over supervised methods. The parameters were chosen in order to give a reasonable separation between the various methods. The dataset is shown in figure (3). For the next dataset we took the digits 0 and 4 from the MNIST training set and randomly sampled 500 examples from each class. We call the resulting dataset reduced MNIST or rMNIST. The third dataset consisted in data from the Oil pipe dataset. This dataset was first considered for machine learning in the work of Bishop and James (1993). The covariates consist in attenuation measurements of gamma rays passed through an oil pipe. In our case we randomly select 500 examples each corresponding to homogeneous and annular flow and the task is to distinguish between these two classes.




Figure 3: The double helix dataset. The red and blue points show the two different classes to be classified.

For each dataset we varied the number of labelled points and also performed 10 randomized training / test partitions. Before comparing the algorithms we ruled out the existence of a trivial semi-supervised solution in terms of a bipartite nearest neighbour graph on the Euclidean distance between covariates.

All supervised algorithms were set up as in section 4.2. The Harmonic functions algorithm was given an isotropic squared exponential weight function. The single length scale hyperparameter was trained using cross validation. The SLGCP again had one free parameter, the common length scale of the covariance function and this was fixed using 10 fold cross validation on the 0-1 loss. The Harmonic functions algorithm and SLGCP were trained ‘transductively’ ie. they had access to the test points as unlabelled data during training. With 10 labelled points, at the optimal length scale, the 10 folds of cross validation for SLGCP took 1 minute and 25 seconds on the same computer described in section 4.2.

Figures (4), (5) and (6) show the results of the experiment. As we would expect the semi-supervised SLGCP generally outperforms its supervised counterpart LGCP and all the other supervised methods, often by a substantial margin. In the majority of cases SLGCP and the Harmonic functions algorithms perform similarly well the only exception being for the smallest number of training points on the double helix dataset where the Harmonic functions algorithm has a lower error rate. A common feature of the datasets is that for each class the covariates lie near a low dimensional manifold in the higher dimensional space, separated by a complex decision surface and the unlabelled data can make a large difference in the quality of prediction by helping to identify these manifolds. Our results show the SLGCP exploiting this advantage.




Figure 4: Test error for the double helix data set as a function of number of training points. The error bars are the standard error of the mean with averaging over 10 random partitions.

## 5 Conclusions

In this paper we extended the LGCP model to semi-supervised tasks and gave the first practical demonstration of the LGCP model on supervised and semi-supervised tasks. The linear time supervised prediction meant that we were able to demonstrate relatively large supervised datasets. Despite its simplicity, there are domains where performance is comparable to well established methods. We gave the first full description of the Markov random field showing that it was pairwise submodular- a necessary property to use graph min-cut methods. We also described intriguing novel connections to classical nonparametric methods.

In terms of further work a promising avenue is to incorporate more complex inductive biases into the mean and covariance functions, exploiting literature on Markov random fields (Koller and Friedman, 2009) for parameter learning. For the semi-supervised learning method, as mentioned in section 3.4 there is scope to combine the existing literature on randomized MAP with these methods (Blum et al., 2004; Hazan et al., 2013). In Bayesian nonparametric terms this method is an example of the largely untapped potential for combining the rich theory of spatial point processes with the important application domain of classification and we believe this is a promising area for further research.

## 6 Acknowledgements

We would like to acknowledge the support of EPSRC grant EP/I036575/1 and a Google Focussed Research Award. We would like to thank Peter McCullagh for comments that improved the presentation and clarity of the paper.

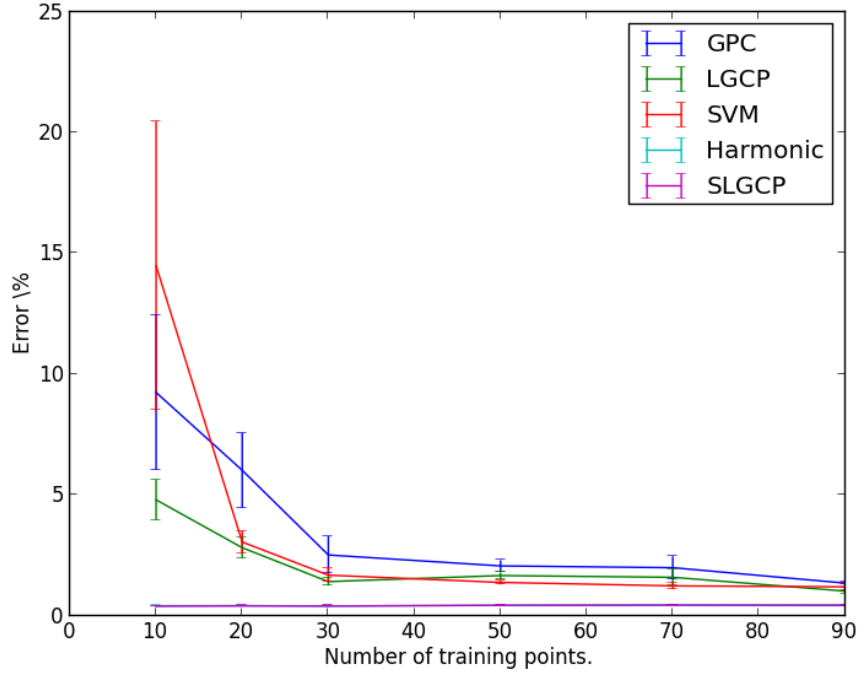


Figure 5: Test error for the rMNIST data set as a function of number of training points. The error bars are the standard error of the mean with averaging over 10 random partitions. The harmonic functions algorithm is indistinguishable from SLGCP (see text).

## Appendix: Some comments on the comparison between LGCP and the per-manental Cox process classifier

In this section we compare classification based on the per-manental Cox process (McCullagh and Møller, 2006) to the LGCP classifier. The comparison has been made before in the existing literature theoretically (McCullagh and Yang, 2006; McCullagh, 2008) and the per-manental process classifier has been studied experimentally (Yang et al., 2012). Here we provide some summary comments on this existing literature and discuss practical insights from our recent work on the LGCP, with a particular bias towards a machine learning perspective.

For the per-manental Cox process both the product density and the Janossy density are available in closed form (McCullagh and Møller, 2006). We define a new function  $\hat{C} = 2C$ . The product density is given by:

$$m_{[K]}(x_1, x_2, \dots, x_K) = \text{Per}_\alpha(\hat{C}[\mathbf{x}]). \quad (36)$$

Here  $\text{Per}_\alpha(B)$  is the alpha permanent of a matrix  $B$ . The local Janossy density with respect to Lebesgue measure on  $(\mathbb{R}^D)^K$  and bounded Borel set  $A$  is given by

$$\theta(x_1, x_2, \dots, x_K | A) = \exp\{-\alpha D\} \text{Per}_\alpha(\tilde{C}[\mathbf{x}]) \quad (37)$$

where  $\tilde{C}$  is defined in terms of the eigenvalues  $\{\lambda_i\}$  and eigenfunctions  $\{\mathbf{e}_i\}$  of  $\tilde{C}$  with respect to Lebesgue measure over  $A$  using the relation:

$$\tilde{C}(x, x') = \sum_{i=1}^{\infty} \frac{\lambda_i}{(1 + \lambda_i)} \mathbf{e}(\mathbf{x}) \mathbf{e}(\mathbf{x}') \quad (38)$$

and  $D = \sum_{i=0}^{\infty} \log(1 + \lambda_i)$ . We may in fact obtain closed form predictive equations for classification using the product density (McCullagh, 2008) or using the Janossy density (Yang et al., 2012). By contrast a closed form is only available for the product density in the LGCP case (Møller et al., 1998).




Figure 6: Test error for the oil pipe data set as a function of number of training points. The error bars are the standard error of the mean with averaging over 10 random partitions. The harmonic functions algorithm is indistinguishable from SLGCP (see text).

Unlike the log-Gaussian Cox process, the family of permanent Cox processes is closed under superposition which is another appealing theoretical property. Further, McCullagh and Yang (2006) are able to extend their classification model to an unbounded number of classes, effectively deriving a kernel generalization of the Chinese restaurant process. It seems difficult to imagine a similar extension for log-Gaussian Cox processes.

The catch however lies in the difference between ‘closed form’ and tractable. The alpha permanent of a matrix is  $\#P$ -Hard to compute (Valiant, 1979) and although approximating the ratio of two related alpha permanents seems to be an easier problem (Kou and McCullagh, 2009), even approximating this ratio for prediction currently requires  $O(N^3)$  (Yang et al., 2012) for the four-cycle approximation as compared to the  $O(N)$  for the LGCP. This means that the log Gaussian variant will be of relatively more interest when scalability is a factor. Further, the connection between the log Gaussian variant and the Potts model (McCullagh, 2008) allows progress to be made on the semi-supervised learning problem and the supervised case has interesting links to classical non-parametric estimators. Thus we believe that both variants merit continuing investigation.

## References

- Ackley, H., Hinton, E., and Sejnowski, J. (1985). A learning algorithm for Boltzmann machines. *Cognitive Science*, pages 147–169.
- Bengio, Y., Delalleau, O., and Le Roux, N. (2005). The curse of dimensionality for local kernel machines. Technical Report 1258, Département d’informatique et recherche opérationnelle, Université de Montréal.
- Besag, J. (1974). Spatial Interaction and the Statistical Analysis of Lattice Systems. *Journal of the Royal Statistical Society. Series B (Methodological)*, 36(2):192–236.
- Bishop, C. and James, G. (1993). Analysis of multiphase flows using dual-energy gamma densitometry and neural

- networks. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 327(23):580 – 593.
- Blum, A. and Chawla, S. (2001). Learning from labeled and unlabeled data using graph mincuts. In *Proceedings of the Eighteenth International Conference on Machine Learning, ICML '01*, pages 19–26, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.
- Blum, A., Lafferty, J., Rwebangira, M. R., and Reddy, R. (2004). Semi-supervised learning using randomized mincuts. In *Proceedings of the Twenty-first International Conference on Machine Learning, ICML '04*, pages 13–, New York, NY, USA. ACM.
- Boykov, Y. and Kolmogorov, V. (2004). An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 26(9):1124–1137.
- Boykov, Y., Veksler, O., and Zabih, R. (2001). Fast approximate energy minimization via graph cuts. *IEEE Trans. Pattern Anal. Mach. Intell.*, 23(11):1222–1239.
- Ciresan, D. C., Meier, U., and Schmidhuber, J. (2012). Multi-column deep neural networks for image classification. *CoRR*, abs/1202.2745.
- Coates, A., Ng, A. Y., and Lee, H. (2011). An analysis of single-layer networks in unsupervised feature learning. In *AISTATS*, pages 215–223.
- Cox, D. (1958). *Planning of experiments*. Wiley series in probability and mathematical statistics: Applied probability and statistics. Wiley.
- Cox, D. R. (1955). Some Statistical Methods Connected with Series of Events. *Journal of the Royal Statistical Society. Series B (Methodological)*, 17(2):129–164.
- Daley, D. J. and Vere-Jones, D. (2003). *An introduction to the theory of point processes. Vol. I. Probability and its Applications* (New York). Springer-Verlag, second edition. Elementary theory and methods.
- Hazan, T., Maji, S., and Jaakkola, T. (2013). On sampling from the gibbs distribution with random maximum a-posteriori perturbations. In Burges, C., Bottou, L., Welling, M., Ghahramani, Z., and Weinberger, K., editors, *Advances in Neural Information Processing Systems 26*, pages 1268–1276. Curran Associates, Inc.
- Jerrum, M. and Sinclair, A. (1993). Polynomial-time approximation algorithms for the Ising model. *SIAM Journal on Computing*, 22:1087–1116.
- Joachims, T. (1999). Advances in kernel methods. chapter Making Large-scale Support Vector Machine Learning Practical, pages 169–184. MIT Press, Cambridge, MA, USA.
- Kingman, J. F. C. (1993). *Poisson processes*. Oxford Studies in Probability. The Clarendon Press Oxford University Press, New York. Oxford Science Publications.
- Koller, D. and Friedman, N. (2009). *Probabilistic Graphical Models: Principles and Techniques*. MIT Press.
- Kolmogorov, V. and Zabih, R. (2004). What energy functions can be minimized via graph cuts? *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26:65–81.
- Kou, S. C. and McCullagh, P. (2009). Approximating the alpha-permanent. *Biometrika*, 96(3):635–644.
- Lafferty, J. D., McCallum, A., and Pereira, F. C. N. (2001). Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of the Eighteenth International Conference on Machine Learning*, pages 282–289, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.
- Lal, T., Schroder, M., Hinterberger, T., Weston, J., Bogdan, M., Birbaumer, N., and Scholkopf, B. (2004). Support vector channel selection in bci. *Biomedical Engineering, IEEE Transactions on*, 51(6):1003–1010.
- LeCun, Y., Bottou, L., Y., B., and P., H. (1998). Gradient-based learning applied to document recognition. In *Proceedings of the IEEE*, pages 2278–2324.



- LeCun, Y., C., C., and C., B. (2014). The MNIST database of handwritten digits. <http://yann.lecun.com/exdb/mnist>.
- McCullagh, P. (2008). Sampling bias and logistic models. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 70(4):643–677.
- McCullagh, P. and Møller, J. (2006). The permanental process. *Adv. Appl. Probab*, 38:873–888.
- McCullagh, P. and Yang, J. (2006). Stochastic classification models. *Proceedings of the International Congress of Mathematicians*.
- Møller, J., Syversveen, A. R., and Waagepetersen, R. P. (1998). Log Gaussian Cox processes. *Scandinavian Journal of Statistics*, 25(3):451–482.
- Møller, J. and Waagepetersen, R. (2004). *Statistical Inference and Simulation for Spatial Point Processes*. Chapman & Hall/CRC Monographs on Statistics & Applied Probability. Taylor & Francis.
- Naish-Guzman, A. and Holden, S. (2008). The generalized FITC approximation. In Platt, J., Koller, D., Singer, Y., and Roweis, S., editors, *Advances in Neural Information Processing Systems 20*, pages 1057–1064. MIT Press, Cambridge, MA.
- Orbanz, P. and Teh, Y. W. (2010). Bayesian nonparametric models. In *Encyclopedia of Machine Learning*. Springer.
- Parzen, E. (1962). On estimation of a probability density function and mode. *The Annals of Mathematical Statistics*, 33(3):pp. 1065–1076.
- Pitman, J. (2006). *Combinatorial stochastic processes*, volume 1875 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin.
- Rasmussen, C. E. and Nickisch, H. (2010). Gaussian processes for machine learning (GPML) toolbox. *Journal of Machine Learning Research*, 11:3011–3015.
- Rasmussen, C. E. and Williams, C. K. I. (2006). *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press.
- Ripley, B. D. (1996). *Pattern Recognition and Neural Networks*. Cambridge University Press.
- Rosenblatt, M. (1956). Remarks on Some Nonparametric Estimates of a Density Function. *The Annals of Mathematical Statistics*, 27(3):832–837.
- Tenenbaum, J. B., Silva, V., and Langford, J. C. (2000). A Global Geometric Framework for Nonlinear Dimensionality Reduction. *Science*, 290(5500):2319–2323.
- Valiant, L. (1979). The complexity of computing the permanent. *Theoretical Computer Science*, 8(2):189 – 201.
- Veksler, O. (1999). *Efficient Graph-based Energy Minimization Methods in Computer Vision*. PhD thesis, Ithaca, NY, USA. AAI9939932.
- Wu, F. Y. (1982). The Potts model. *Reviews of Modern Physics*, 54(1):235–268.
- Yang, J., Miescke, K., and McCullagh, P. (2012). Classification based on a permanental process with cyclic approximation. *Biometrika*, 99(4):775–786.
- Zhang, Y., Brady, M., and Smith, S. (2001). Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm. *Medical Imaging, IEEE Transactions on*, 20(1):45–57.
- Zhu, X., Ghahramani, Z., and Lafferty, J. (2003). Semi-supervised learning using Gaussian fields and harmonic functions. In *ICML*, pages 912–919.